



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/332,264	06/11/1999	THOMAS HUNTINGTON WOOD	WOOD27/56115	2709
26291	7590	12/28/2004	EXAMINER	
MOSER, PATTERSON & SHERIDAN L.L.P. 595 SHREWSBURY AVE, STE 100 FIRST FLOOR SHREWSBURY, NJ 07702			LI, SHI K	
			ART UNIT	PAPER NUMBER
			2633	

DATE MAILED: 12/28/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

MAILED  
DEC 28 2004  
GROUP 2600



UNITED STATES PATENT AND TRADEMARK OFFICE

---

COMMISSIONER FOR PATENTS  
UNITED STATES PATENT AND TRADEMARK OFFICE  
P.O. Box 1450  
ALEXANDRIA, VA 22313-1450  
[www.uspto.gov](http://www.uspto.gov)

MAILED

DEC 28 2004

GROUP 2600

**BEFORE THE BOARD OF PATENT APPEALS  
AND INTERFERENCES**

Application Number: 09/332,264

Filing Date: June 11, 1999

Appellant(s): WOOD, THOMAS HUNTINGTON

---

Joseph Pagnotta  
For Appellant

**EXAMINER'S ANSWER**

This is in response to the appeal brief filed 31 August 2004.

**(1) *Real Party in Interest***

A statement identifying the real party in interest is contained in the brief.

**(2) *Related Appeals and Interferences***

A statement identifying the related appeals and interferences which will directly affect or be directly affected by or have a bearing on the decision in the pending appeal is contained in the brief.

**(3) *Status of Claims***

The statement of the status of the claims contained in the brief is correct.

**(4) *Status of Amendments After Final***

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

**(5) *Summary of Invention***

The summary of invention contained in the brief is correct.

**(6) *Issues***

The appellant's statement of the issues in the brief is correct.

**(7) *Grouping of Claims***

Appellant's brief includes a statement that claims 1, 2, 4, 5 and 7-15 do not stand or fall together and provides reasons as set forth in 37 CFR 1.192(c)(7) and (c)(8).

**(8) *Claims Appealed***

The copy of the appealed claims contained in the Appendix to the brief is correct.

**(9) *Prior Art of Record***

Art Unit: 2633

5,311,344	BOHN	05-1995
5,550,666	ZIRNGIBL	08-1996
6,137,607	FELDMAN	10-2000
6,542,722	SORRELLS	04-2003
6,587,476	LEWIN	07-2003

**(10) Grounds of Rejection**

The following ground(s) of rejection are applicable to the appealed claims:

**Claims 1-2, 4-5, 7, 10 and 12-15** are rejected under 35 U.S.C. 103(a) as being unpatentable over Bohn et al. (U.S. Patent 5,311,344) in view of Lewin et al. (U.S. Patent 6,587,476 B1).

Bohn et al. discloses in FIG. 1 a data communication system comprising a head-end 2, a splitter 3, a first network unit 5<sub>1</sub> and a second network unit 5<sub>2</sub>. The difference between Bohn et al. and the claimed invention is that (a) Bohn et al. does not include an Ethernet adapter circuit in the head-end; and (b) Bohn et al. does not teach the Ethernet interface for providing the upstream data.

Ethernet is a popular network interface and can be found in most computers for interconnecting with other computers. Lewin et al. emphasizes the fact in "Background of the Invention" Section and teaches in FIG. 1 the use of 10 BaseT Ethernet interfaces for receiving data from subscribers. Lewin et al. also teaches in FIG. 7 the use of an Ethernet switch to combine the data from individual subscribers. One of ordinary skill in the art would have been motivated to combine the teaching of Lewin et al. with the data communication system of Bohn et al. to use Ethernet interface for receiving subscriber data because Ethernet interfaces are

Art Unit: 2633

popularly found in home and office computers, and to include a Ethernet switch in the head-end for combining data from individual subscribers because the Ethernet switch provides high bandwidth duplex data communication for each subscriber and allows the interconnection to other networks. In using Ethernet switch, there is no collision between the different ports. Thus it would have been obvious to one of ordinary skill in the art at the time the invention was made to use Ethernet interface for receiving upstream data and include a Ethernet switch in the head-end for combining data from individual subscribers and interconnecting to other networks, as taught by Lewin et al., in the data communication system of Bohn et al. because Ethernet interfaces are popularly equipped in most home and office computers.

Regarding claim 2, Bohn et al. uses subcarrier technology as described in col. 4, lines 39-65.

Regarding claim 4, Bohn et al. uses optical fibers for connecting the network units and the splitter. Therefore the network units are optical network units.

Regarding claim 5, the modified data communication system of Bohn et al. and Lewin et al. would use Ethernet data format, which is a packet format, for the data.

Regarding claims 7 and 12, the modified data communication system of Bohn et al. and Lewin et al. would include an Ethernet adapter, a modulator (VCO 57 of FIG. 2 of Bohn et al.) and a transmitter (laser 55 of FIG. 2 of Bohn et al.).

Regarding claim 10, Bohn et al. suggests the use of FSK in col. 4, line 61.

Regarding claim 13, Bohn et al. includes in FIG. 1 downstream data to the network units.

Regarding claim 14, Bohn et al. suggests in FIG. 1 that the head-end includes a transmitter 23, a receiver 24 and a wavelength-division multiplexing device (coupler 21). Bohn

Art Unit: 2633

et al. also suggests in FIG. 2 that each network unit includes a transmitter 55, a receiver 52 and a wavelength-division multiplexing device 51.

Regarding claim 15, Bohn et al. suggests in FIG. 1 that the receiver in the head-end and the transmitter in the network units operate at 1.5  $\mu\text{m}$  and the transmitter in the head-end and the receivers in the network units operate at 1.3  $\mu\text{m}$ .

**Claim 8** is rejected under 35 U.S.C. 103(a) as being unpatentable over Bohn et al. and Lewin et al. as applied to claims 1-2, 4-5, 7, 10 and 12-15 above, and further in view of Feldman (U.S. Patent 6,137,607).

Bohn et al. and Lewin et al. have been discussed above in regard to claims 1-2, 4-5, 7, 10 and 12-15. The difference between the modified data communication system of Bohn et al. and Lewin et al. and the claimed invention is that Bohn et al. and Lewin et al. do not include a bias control circuit. Feldman et al. describes the operation of the bias control 204 in col. 2, lines 60-67 such that the bias control circuit shuts off the laser (transmitter) in the absence of user data. One of ordinary skill in the art would have been motivated to combine the teaching of Feldman et al. with the modified data communication system of Bohn et al. and Lewin et al. because the bias control circuit eliminates optical beat interference (OBI) as described in col. 2, lines 60-61 of Feldman. Thus it would have been obvious to one of ordinary skill in the art at the time the invention was made to incorporate the bias control circuit, as taught by Feldman et al., into the modified system of Bohn et al. and Lewin et al. to eliminate optical beat interference.

Art Unit: 2633

**Claim 9** is rejected under 35 U.S.C. 103(a) as being unpatentable over Bohn et al. and Lewin et al. as applied to claims 1-2 above, and further in view of Sorrells et al. (U.S. Patent 6,542,722 B1).

Bohn et al. and Lewin et al. have been discussed above in regard to claims 1-2, 4-5, 7, 10 and 12-15. The difference between the modified data communication system of Bohn et al. and Lewin et al. and the claimed invention is the modulation method for upstream data. Bohn et al. suggests the use of FSK while the claimed invention uses QPSK. Sorrells et al. teaches in col. 11, lines 49-60 techniques for modulation. These different techniques for modulation are considered as equivalents and the choice of one technique over the others depends on the particular application, e.g., the number of subscribers, data rate and cost. Where the claimed differences involve the substitution of interchangeable or replaceable equivalents and the reason for the selection of one equivalent for another was not to solve an existent problem, such substitution has been judicially determined to have been obvious. See *In re Ruff*, 118, USPQ 343 (CCPA 1958). Thus it would have been obvious to one of ordinary skill in the art at the time the invention was made to use QPSK as a modulation method, as taught by Sorrells et al., in the modified system of Bohn et al. and Lewin et al. as a design choice based on the particular application.

**Claim 11** is rejected under 35 U.S.C. 103(a) as being unpatentable over Bohn et al. and Lewin et al. as applied to claims 1-2 above, and further in view of Zirngibl (U.S. Patent 5,550,666).

Bohn et al. and Lewin et al. have been discussed above in regard to claims 1-2, 4-5, 7, 10 and 12-15. The difference between the modified data communication system of Bohn et al. and Lewin et al. and the claimed invention is the wavelength for the upstream data. Zirngibl teaches in col. 2, line 39-41 the use of 1.3  $\mu\text{m}$  wavelength for upstream data. Certain optical fiber has a minimal absorption loss around 1.3  $\mu\text{m}$ . Thus it would have been obvious to one of ordinary skill in the art at the time the invention was made to use 1.3  $\mu\text{m}$  wavelength for upstream data, as taught by Zirngibl, in the modified system of Bohn et al. and Lewin et al. because certain optical fiber has a minimal absorption loss at wavelength around 1.3  $\mu\text{m}$ .

**(11) *Response to Argument***

In page 9 of the Brief, Appellants argue that "mere popularity of a given feature in the references cannot be used as a suggestion or a motivation to combine the teachings because the references must suggest the desirability of a claimed invention. The Examiner recognizes that obvious can only be established by combining or modifying the teachings of the prior art to produce the claimed invention where there is some teaching, suggestion, or motivation to do so found either in the references themselves or in the knowledge generally available to one of ordinary skill in the art. See *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed Cir. 1988) and *In re Jones*, 958 F.2d 347, 21 USPQ2d 1941 (Fed. Cir. 1992). In this case, popularity of Ethernet is well known in the art and considered as knowledge generally available to one of ordinary skill in the art. In fact, the popularity of Ethernet is also taught or suggested by Lewin et al. in the Background of the Invention section. Therefore, the argument cannot sustain.

In page 10 of the Brief, Appellants state "The level of skill in the art cannot be relied upon to provide the suggestion to combine references" and cites *Site Corp v. VSI International*



Art Unit: 2633

Inc. 50 USPQ 1161 (Fed. Cir. 1999). The Appellants misquote or misunderstand the law case. Section 7 of the case states "Infringement defendant cannot rely on what it presumes is level of knowledge of one of ordinary skill in the art at time of invention to supply suggestion to combine prior art references" (emphasis added). It then states "there is substantial evidence to support conclusion that one of ordinary skill in art would not have known to make asserted combination in present case" (emphasis added). In the Site Corp v. VSI International Inc. case, VSI suggests to bridge the gap between prior art and claimed invention with "presumed" level of knowledge of one of ordinary skill in the art. VSI is unable to point to any specific teaching or suggestion for making this combination. (See page 1171 of 50 USPQ 1161.) In the instant rejection, the Examiner relies on teaching of Lewin et al. to bridge the gap between Bohn et al. and the claimed invention. To prove that popularity of Ethernet is in fact well known in the art, the Examiner attaches Frazier (H. Frazier, "The 802.3z Gigabit Ethernet Standard", IEEE Network, May/June 1998) where it describes Ethernet as "the world's most popular" standard in the first sentence. Therefore, the popularity of Ethernet is level of knowledge of one of ordinary skill in the art, and the combination of Bohn et al. and Lewin et al. is well motivated and obvious.

In page 11 of the Brief, Appellants argue that "mere popularity of a given feature cannot be used as suggestion or motivation to combine the teachings". Here, Appellants give a conclusive statement to rule out certain suggestion or motivation and cannot be accepted. Popularity of a particular interface raises compatibility issue. New product are usually designed to be backward compatible with exist equipment. For example, Fawaz et al. (U.S. Patent 6,654,374 B1) teaches in col. 13, line 65-col. 14 line 3 that "a system in accordance with the invention is not only high performance, but also economical, simple, and flexible, while

Art Unit: 2633

remaining compatible with much of the hardware and software already in place, particularly hardware and software for routers and ethernet interfaces". One of ordinary skill in the art would have been motivated to incorporate popular interface into an apparatus because it makes the apparatus compatible with existing equipment.

Appellants argue that "Bohn achieves collisionless transmission in a completely different manner" and conclude that "there is no advantage to be gained by introducing Ethernet in Bohn when Bohn already solves his problem through the multiplexing technologies". Again, Appellant refuses to recognize popularity, or compatibility with existing equipment, as an advantage and make an assumption that "achieving collisionless transmission" were the motivation for combining the references. Such argument is pointless.

In page 13, Appellants argue that Lewin converts Ethernet frames to VDSL for operating along a twisted pair medium and not optically-based networks. The test for obviousness is not whether the features of the reference may be bodily incorporated into the other to produce the claimed subject matter but simply what the references make obvious to one of ordinary skill in the art. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981). In this case, Bohn teaches an optical subscriber network with head-end and subscriber terminals. Lewin suggests to use an Ethernet interface, which is commonly found in PC, for interfacing subscriber terminals to user equipment (PC) and an Ethernet switch to route Ethernet traffic in the head-end. One of ordinary skill in the art would have been motivated to combine the teaching of Lewin et al. with the data communication system of Bohn et al. to use Ethernet interface for receiving subscriber data because Ethernet interfaces are popularly found in home and office computers.

Art Unit: 2633

Appellants argue that claims 2, 4, 5, 7, 10 and 12-15 are allowable because of their dependence on claim 1. The argument is moot in view of the unpatentability of claim 1 as explained above.

Appellants argue that claims 8, 9 and 11 are allowable because of their dependence on claim 1. The argument is moot in view of the unpatentability of claim 1 as explained above.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

skl

December 23, 2004

Conferees

M. R. Sedighian

Daniel J. Wu

*M. R. Sedighian*

*Daniel J. Wu*

BAKER & MCKENZIE  
805 THIRD AVENUE  
NEW YORK, NY 10022

## The 802.3z Gigabit Ethernet Standard

By Howard Frazier

*This column considers standards development in the IEEE 802.3z Gigabit Ethernet Task Force, part of the IEEE 802.3 CSMA/CD Working Group. IEEE Std 802.3z-1998 was formally approved by the IEEE Standards Board on June 25th, 1998. The author is responsible for the leadership of the Gigabit Project and serves as the Task Force chair of IEEE 802.3z.*

*These perspectives, which will appear in each issue of IEEE*


*Network, are aimed to provide the reader with highlights of IEEE 802 activities to enable better dissemination of these standards into marketable products as well as seek new ideas to be brought into the IEEE 802 arena. I hope the readers find this column useful and look forward to providing future installments. Please send suggestions, criticisms, and requests for future articles to jcarlo@ti.com, IEEE 802 chair.*

The standard for gigabit Ethernet, IEEE Std 802.3z, extends the operating speed of the world's most popular local area network to 1 billion bits per second (1000 Mb/s) for interconnecting high-performance switches, routers, and servers in the backbone of local area networks. Maintaining backward compatibility with the over-100-million-node installed base of 10 Mb/s and 100 Mb/s was a key requirement. Throughout the two-and-a-half-year standard development cycle, the IEEE P802.3z Gigabit Task Force was guided by the overall objective of increasing the data rate of Ethernet technology by embracing the advancing state of the art in network design and physical-layer signaling technology. Soon to be published as a supplement to ANSI/IEEE Std 802.3 (fifth edition), IEEE Std 802.3z represents the next logical step in the evolution of Ethernet local area networks.

As a supplement to ANSI/IEEE Std 802.3, the standard for Gigabit Ethernet had a solid base to build on. IEEE Std 802.3z consists of a series of updates to the 802.3 base standard plus a group of new clauses (chapters) which describe the characteristics unique to 1000 Mb/s operation. The standards produced by the IEEE 802.3 Working Group address the two lowest layers of the International Organization for Standardization (ISO) seven-layer reference model, the data link and physical layers. Within these two layers, the specification is divided into sublayers, separated by well defined interfaces. Figure 1 presents a block diagram identifying the various components of IEEE Std 802.3z.

Starting from the top of the diagram, the media access control (MAC) sublayer describes the algorithms used to control the transmission and reception of packets on an Ethernet network. IEEE Std 802.3z encompasses both the newer full-duplex MAC and the classic carrier sense multiple access with collision detection (CSMA/CD) MAC. The full-duplex MAC is described in IEEE Std 802.3x, and allows point-to-point simultaneous transmit and receive communications with a dedicated full data rate per station. The classic Ethernet CSMA/CD is based on a shared Ethernet where all stations share the same data rate. While most implementations will take advantage of the contention-free access and flexible topologies permitted by full-duplex operation, the P802.3z specification also extends the CSMA/CD MAC to work at 1000 Mb/s. A technique known as *carrier extension* was added to the CSMA/CD MAC to overcome the inherent limitation of the collision detection algorithm which mandated that the round-trip delay between any two stations could not exceed the time required to transmit the smallest allowable frame. An optional feature called *frame bursting* was defined to improve the throughput of gigabit CSMA/CD LANs.

When operating in full-duplex mode, the MAC relies on the fact that the underlying serial communications link inherently supports simultaneous transmission and reception without interference between the transmitted and received signals. Essentially, the CSMA/CD protocol is disabled, and since there is no round-trip delay constraint, carrier extension is not



■ Figure 1. Gigabit Ethernet Layer diagram.

necessary. Furthermore, the span of any given link may be as long as the underlying physical layer permits. Full-duplex operation was introduced into IEEE Std 802.3 by the P802.3x project, which also specified a mechanism for link-level flow control. In the event of transient congestion on a link, a receiver can inhibit the transmission of frames by emitting a Pause frame, which instructs the transmitter to withhold further transmissions until some amount of time transpires.

The optional gigabit media-independent interface (GMII) definition allows MAC and physical-layer (PHY) implementations from different vendors to interoperate, and provides a means by which the future 1000BASE-T PHY can be attached to the MAC. The GMII also provides a convenient division between the high-density digital logic functions associated with the MAC, and the high-speed mixed signal functions associated with the PHY. This partition allows appropriate design and fabrication processes to be applied to each of these functions.

The GMII delivers 8-bit octets to the physical coding sublayer (PCS) on the transmit path, and accepts 8-bit octets from the PCS on the receive path at a rate of 125 million octets (1 billion bits) per second. The 1000BASE-X PCS borrows heavily from the NCITS T11 Fibre Channel standard, using the same 8B/10B code set. The P802.3z Task Force adapted the Fibre Channel code to gigabit Ethernet after stripping the coding down to the bare essentials. Starting with a mature and well understood code will save vendors time in the product development process and also speed the standards process.

In a fashion similar to the 100BASE-X PCS defined in IEEE Std 802.3u, the 1000BASE-X PCS uses a continuous signaling scheme, wherein an active signal is always present on the medium. The idle period between packets is filled with a special pattern of symbols which contain a sufficient number of transitions between the logic 1 and logic 0 states so that a receiver's phase-locked loop can maintain clock recovery on a continuous basis. Another special sequence of symbols is used to delimit the beginning and end of a packet.

The 10-bit symbols produced by the PCS are serialized by

the physical medium attachment (PMA) sublayer. Commercially available Serializer/Deserializer (SerDes) components first developed for Fibre Channel can be used for gigabit Ethernet. Only modest additional testing requirements are needed, mainly due to the fact that gigabit Ethernet operates at a clock rate of 1.25 Gb/s (1 Gb/s data rate) while Fibre Channel operates at a clock rate of 1 Gb/s (800 Mb/s data rate). This reuse of off-the-shelf components has contributed to the rapid development and deployment of gigabit Ethernet systems. The PMA sublayer is also responsible for recovering a clock reference from the received data stream. The expansion which results from encoding 8-bit octets into 10-bit symbols requires a signaling rate of 1.25 Gbaud at the serial interface to the medium.

The PCS includes a function referred to as *auto negotiation*, which is a link startup and initialization procedure first defined in IEEE Std 802.3u for 100BASE-T. Within IEEE Std 802.3z, auto negotiation is used to select between the CSMA/CD and full-duplex operating modes, and to select whether the Pause flow control mechanism is enabled or disabled on a link-by-link basis.

At the bottom of the diagram are the 1000BASE-SX and 1000BASE-LX fiber optic transceiver specifications. The 1000BASE-SX specification for short-wavelength laser transceivers supports multi-mode fiber optic links at distances up to 275 m using 62.5  $\mu$  fiber, and 550 m using 50  $\mu$  fiber. Once again, Fibre Channel provided the starting point for the 1000BASE-SX specification, but in addition, the 1000BASE-SX specification embraces VCSELs as well as the older CD style of laser developed for the Fibre Channel market.

1000BASE-LX supports longer distances using higher-cost components, spanning 550 m on 62.5  $\mu$  or 50  $\mu$  fiber, and up to 5 km on single-mode fiber. The 1000BASE-LX laser transmitter is optimized for single-mode fiber, and requires a mode-conditioning patch cord to support multimode fiber optic cable. The patch cord mitigates an effect known as differential mode delay by altering the launch characteristics of a 1000BASE-LX laser transmitter so that the resulting optical beam more closely resembles the overfilled launch pattern produced by a light-emitting diode (LED). Multimode fiber modal bandwidth is specified under the launch conditions produced by an LED. Through use of the mode-conditioning patch cord, a 1000BASE-LX link can be characterized on the basis of the rated modal bandwidth of a multimode fiber. Both 1000BASE-SX and 1000BASE-LX specify the familiar duplex SC optical connector, eliminating the most common installation problem encountered in fiber optic networks, the misconnection of the transmitting and receiving fibers.

IEEE Std 802.3z also includes a specification for a transceiver technology referred to as 1000BASE-CX, which supports shielded copper cables links spanning 25 m. The SerDes component which makes up the PMA sublayer is designed to drive this cable directly, which makes 1000BASE-CX an economically attractive choice for short-distance interconnections, for instance, between devices located within the same rack or within a computer room or telephone closet.

A new project within the IEEE 802.3 Working Group, referred to as 1000BASE-T, is chartered to develop a PHY specification which will support 1000 Mb/s operation on four pairs of category 5 UTP cabling, at a maximum link distance of 100 m. This project is being conducted in the P802.3ab Task Force. 1000BASE-T will take advantage of recent advances in silicon process technology which permit complex high-speed digital signal processing algorithms to be implemented cost effectively. A 1000BASE-T PHY will transmit its signal on all four pairs of wire simultaneously, thus reducing the data rate on each pair to 250 Mb/s. The use of a five-level pulse amplitude modulation scheme further reduces the signaling rate on each pair. Hybrids and digital echo cancellation are used to achieve full-duplex communication.

Network administrators will actively embrace gigabit Ethernet because the familiar management attributes that have been in use in 10 and 100 Mb/s Ethernet networks have been retained. The attribute definitions have been updated to reflect the fact that statistics counters tick 10 times faster for gigabit Ethernet, and certain attributes have been extended to embrace the new physical layers, but the "look and feel" of Ethernet from the network manager's point of view has been preserved.

#### Biography

HOWARD FRAZIER (hfrazier@cisco.com) is employed by Cisco Systems, Inc. within the Workgroup Business Unit. He is chair of the IEEE 802.3z Gigabit Task Force, which developed the Gigabit Ethernet standard. Previously, he was chair of the IEEE 802.3u 100BASE-T Task Force, which developed the standard for Fast Ethernet. Prior to joining Cisco he was employed by Sun Microsystems, Inc. He graduated from Carnegie-Mellon University with a B.S.E.E. in 1983.

#### ATM FORUM/from page 7

proposed which uses the Chernoff method (which in turn uses the Markov, or Chebyshev, inequality) and assumes that local delays at switches are independent and gamma distributed [1].

Other brief discussions centered around the issues at a VC-to-VP aggregation point. One issue was how to determine the QoS of the VP relative to the QoS of the VCs. It has been documented that no general solution is known to this problem. Another issue was how VP sources should handle explicit forward congestion indication (EFCI); the problem is that source behavior 12 requires sources to reset the EFCI state. At a VC-VP boundary, this means that EFCI information from the VC control loop would be lost. The solution proposed for this is to reflect the EFCI state of the incoming data cells back to the VC source, stopping short of a full virtual source/virtual destination implementation.

A number of joint sessions have been held with network management, SAA, and RBB groups. The work with network management centered on how to count valid and invalid RM cells, and how to log traffic descriptors. The work with SAA centered around ABR API issues to provide an interface to query

and set ABR parameters. Another issue was how to request SCR and MBS parameters given the mean and PCR of a video stream. The work with RBB centered around simplification of traffic parameters for residential users, and issues of shared access over asymmetric links such as cable.

Several performance contributions on GFR, IP over ATM, and VS/VD were presented in recent meetings.

#### References

- [1] J. Kenney, Ed., "ATM Traffic Management Living List," ATM Forum LTD-TM-01.07, Apr. 1998.
- [2] ATM Forum Traffic Management, *The ATM Forum Traffic Management Specification Version 4.0*, Apr. 1996; available as [ftp://ftp.atmforum.com/pub/approved-specs/at-tm-0056.000.ps](http://ftp.atmforum.com/pub/approved-specs/at-tm-0056.000.ps).

#### Biography

SHIVKUMAR KALYANARAMAN (shivkuma@ecse.rpi.edu) is an assistant professor at the Department of Electrical, Computer and Systems Engineering at Rensselaer Polytechnic Institute in Troy, NY. He received a B.Tech degree from the Indian Institute of Technology, Madras, India in July 1993, followed by M.S. and Ph.D. degrees in Computer and Information Sciences at the Ohio State University in 1994 and 1997, respectively. His research interests include multimedia networking, traffic management in ATM networks and Internet, Internet pricing, and performance analysis of distributed systems. He is a co-inventor in two patents (the ERICA and OSU schemes for ATM traffic management), and has co-authored several papers and ATM forum contributions in the field of ATM traffic management. He is a member of IEEE-CS and ACM. His homepage is: <http://www.ecse.rpi.edu/Hompages/shivkuma>